

A world map with a dark blue background. Overlaid on the map is a heatmap representing network traffic density. The colors range from dark blue (low density) to yellow and orange (high density). High-density areas are visible in North America, Europe, and parts of Asia. The text 'Dropbox traffic infrastructure' is overlaid on the right side of the map in a large, white, sans-serif font.

# Dropbox traffic infrastructure

Oleg Guba  
SRE, Traffic-team  
[oleg@dropbox.com](mailto:oleg@dropbox.com)

# Dropbox scale

Half a billion of users

180+ countries

Exabytes of users data

Petabytes of metadata

Millions requests per second

Terabits of traffic


# Our Edge network




# Our Edge network



# Why latency matters?


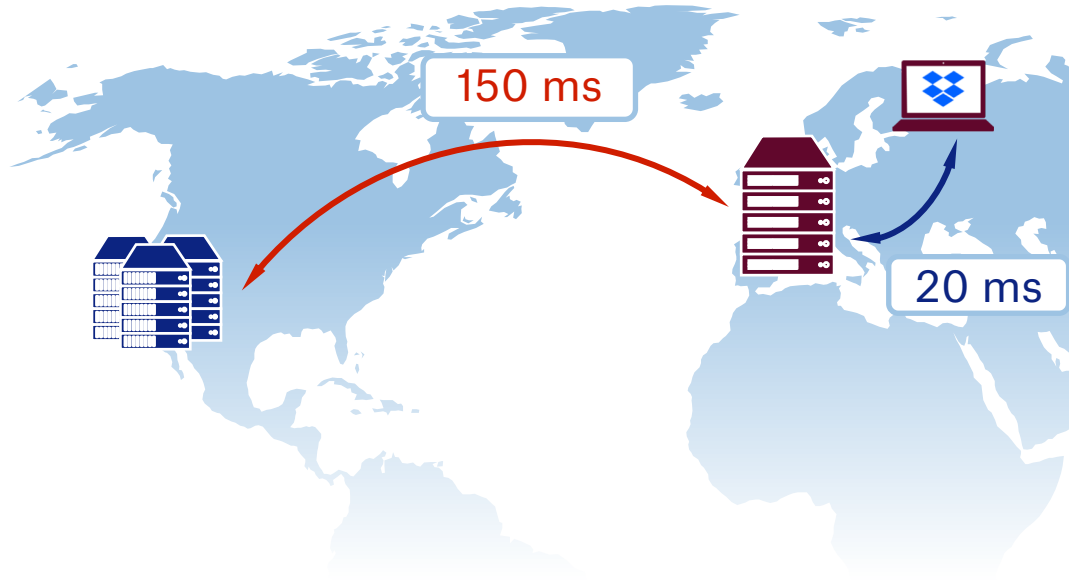


# Why latency matters?




$$4\text{RTT} + \text{server time} = 700\text{ms}$$

# Why latency matters?



$$4\text{RTT} + 1\text{RTT} + \text{server time} = 330\text{ms}$$

# Why latency matters?






# Why we want to be close to users

**Reduce latency for interactive traffic**

**Increase throughput for bulk traffic**

# Global server load balancing (GSLB)




# GSLB techniques

**BGP Anycast**


**DNS**

**URL-based**

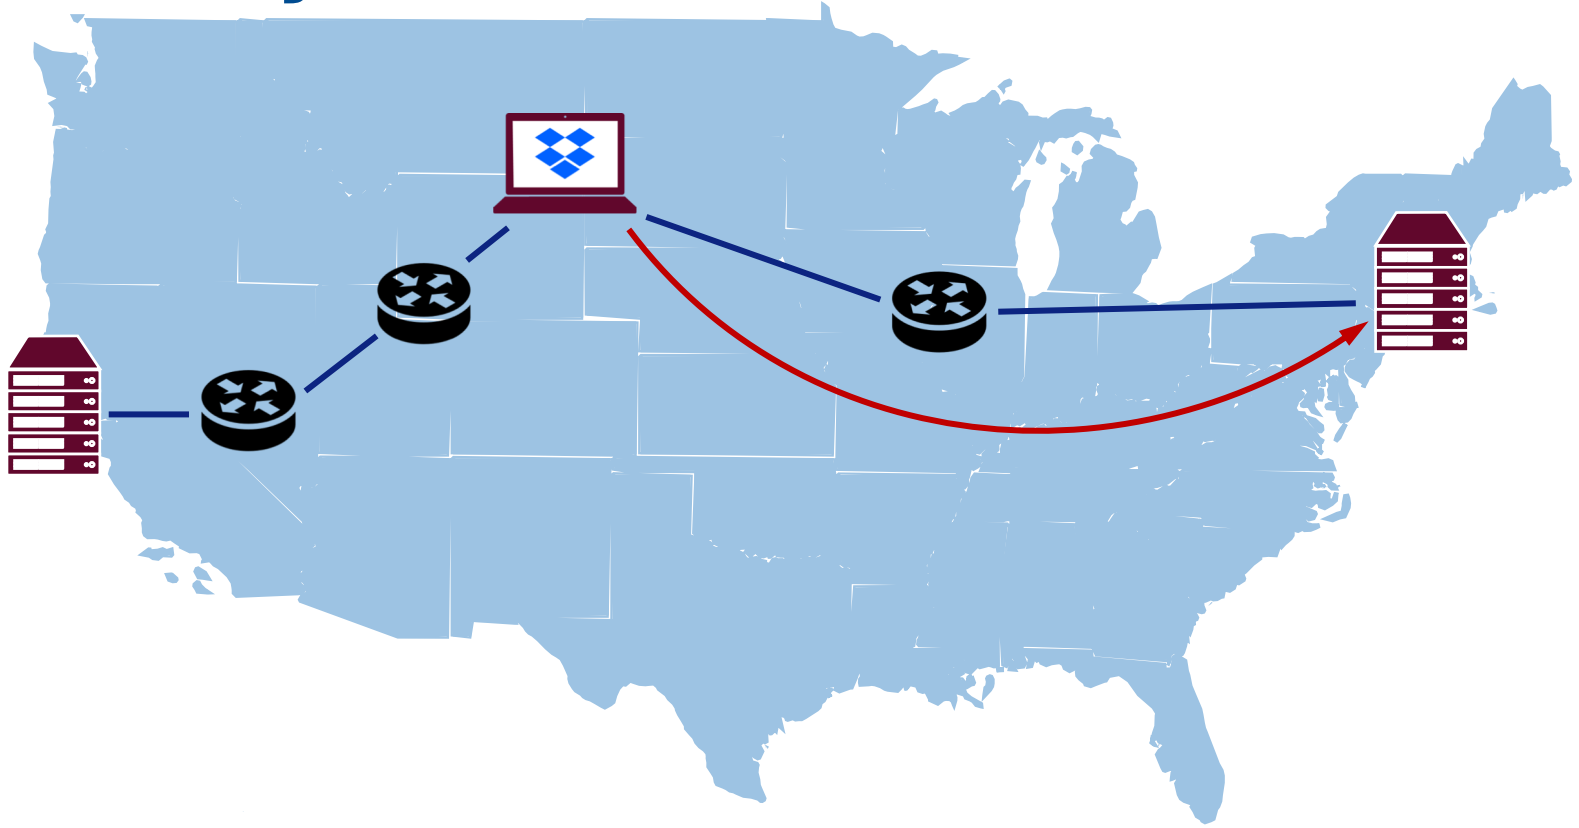
# GSLB: Anycast




# GSLB: Anycast



# GSLB: Anycast issues



# GSLB: Anycast performance



# GSLB: Anycast




Simple  
Automatic failover




BGP is not latency aware  
Almost no traffic control  
No graceful draining




# GSLB: Geo-DNS




# GSLB: Geo-DNS



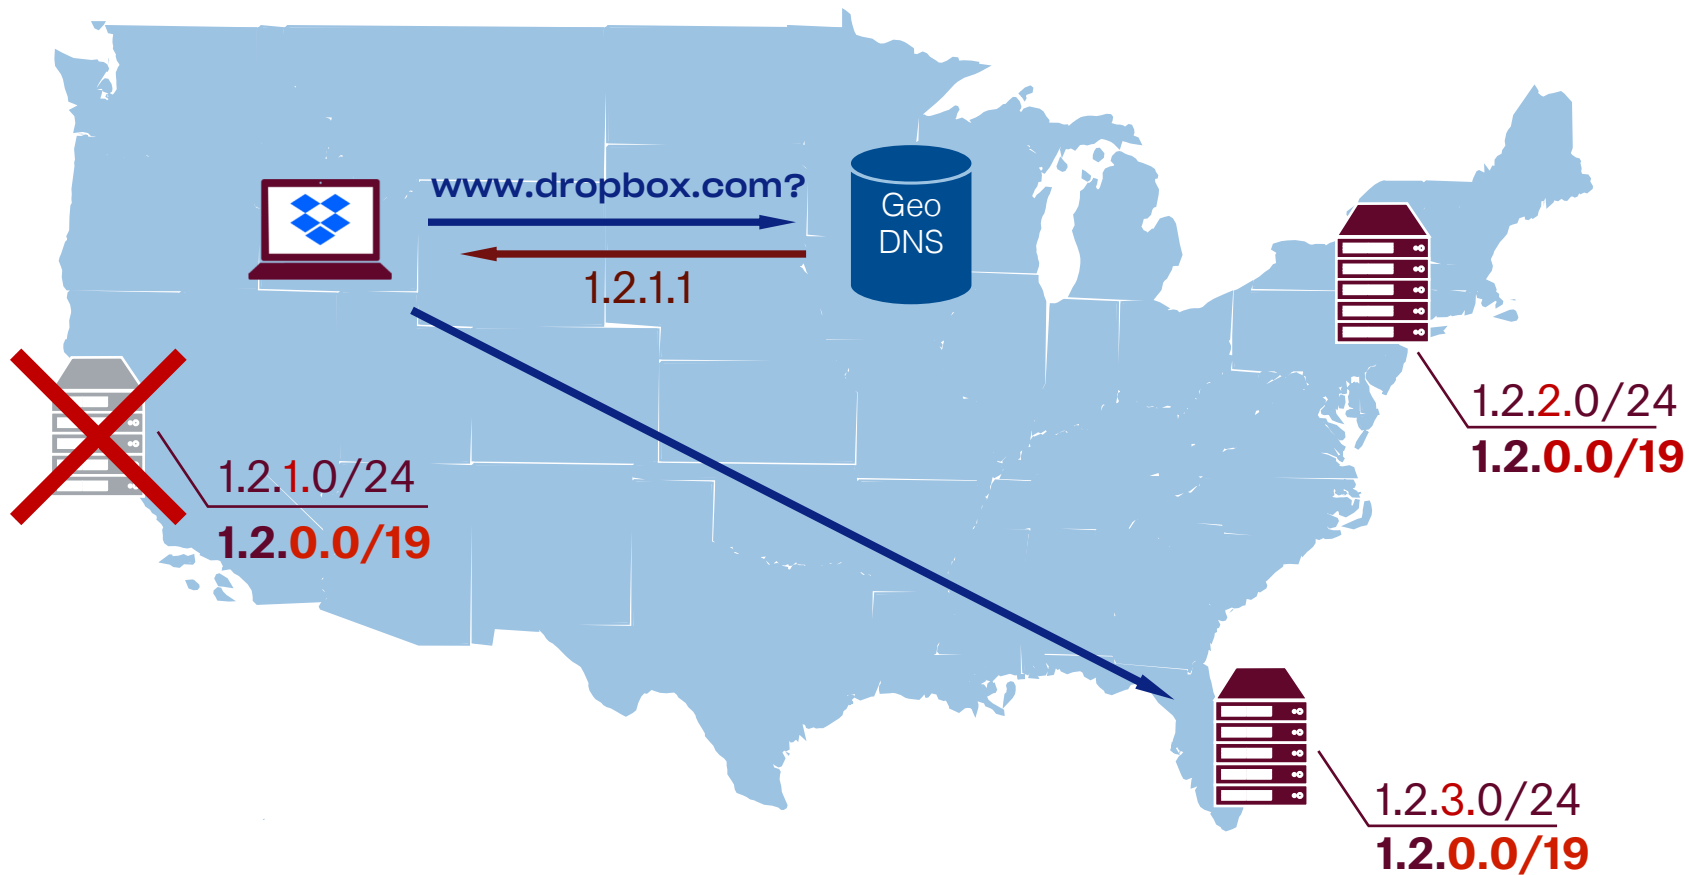
# Geo-DNS: TTL is a lie



# GSLB: Geo-DNS



# GSLB: Geo-DNS



# GSLB: Geo-DNS



**Better routing decision than anycast**

**Traffic control**

**Graceful draining**




**Geo-DNS is not latency aware**

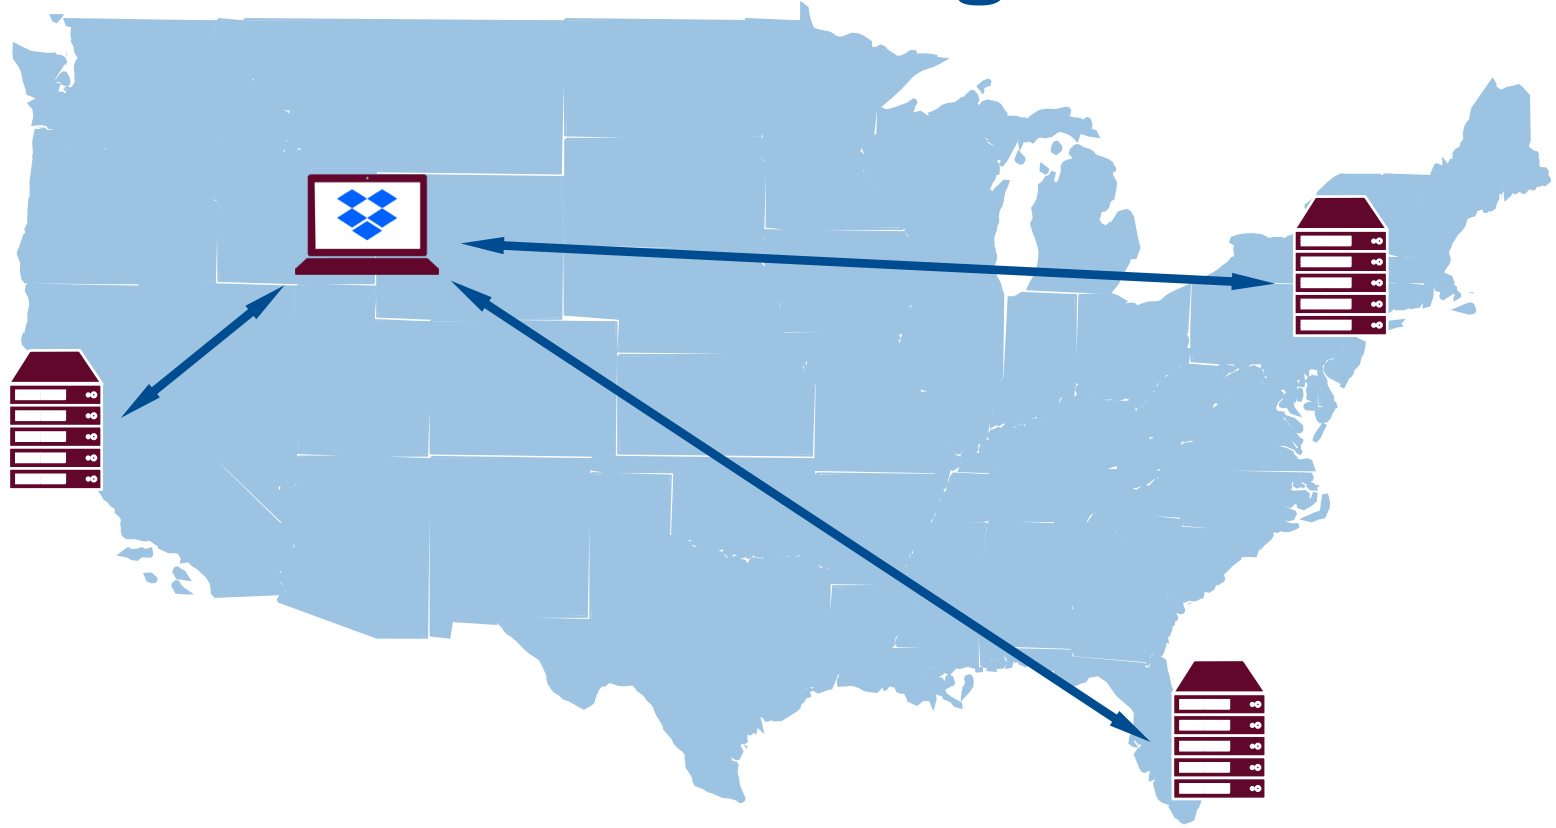
**Precise GeoIP database required**

# GSLB: Real user monitoring (RUM)

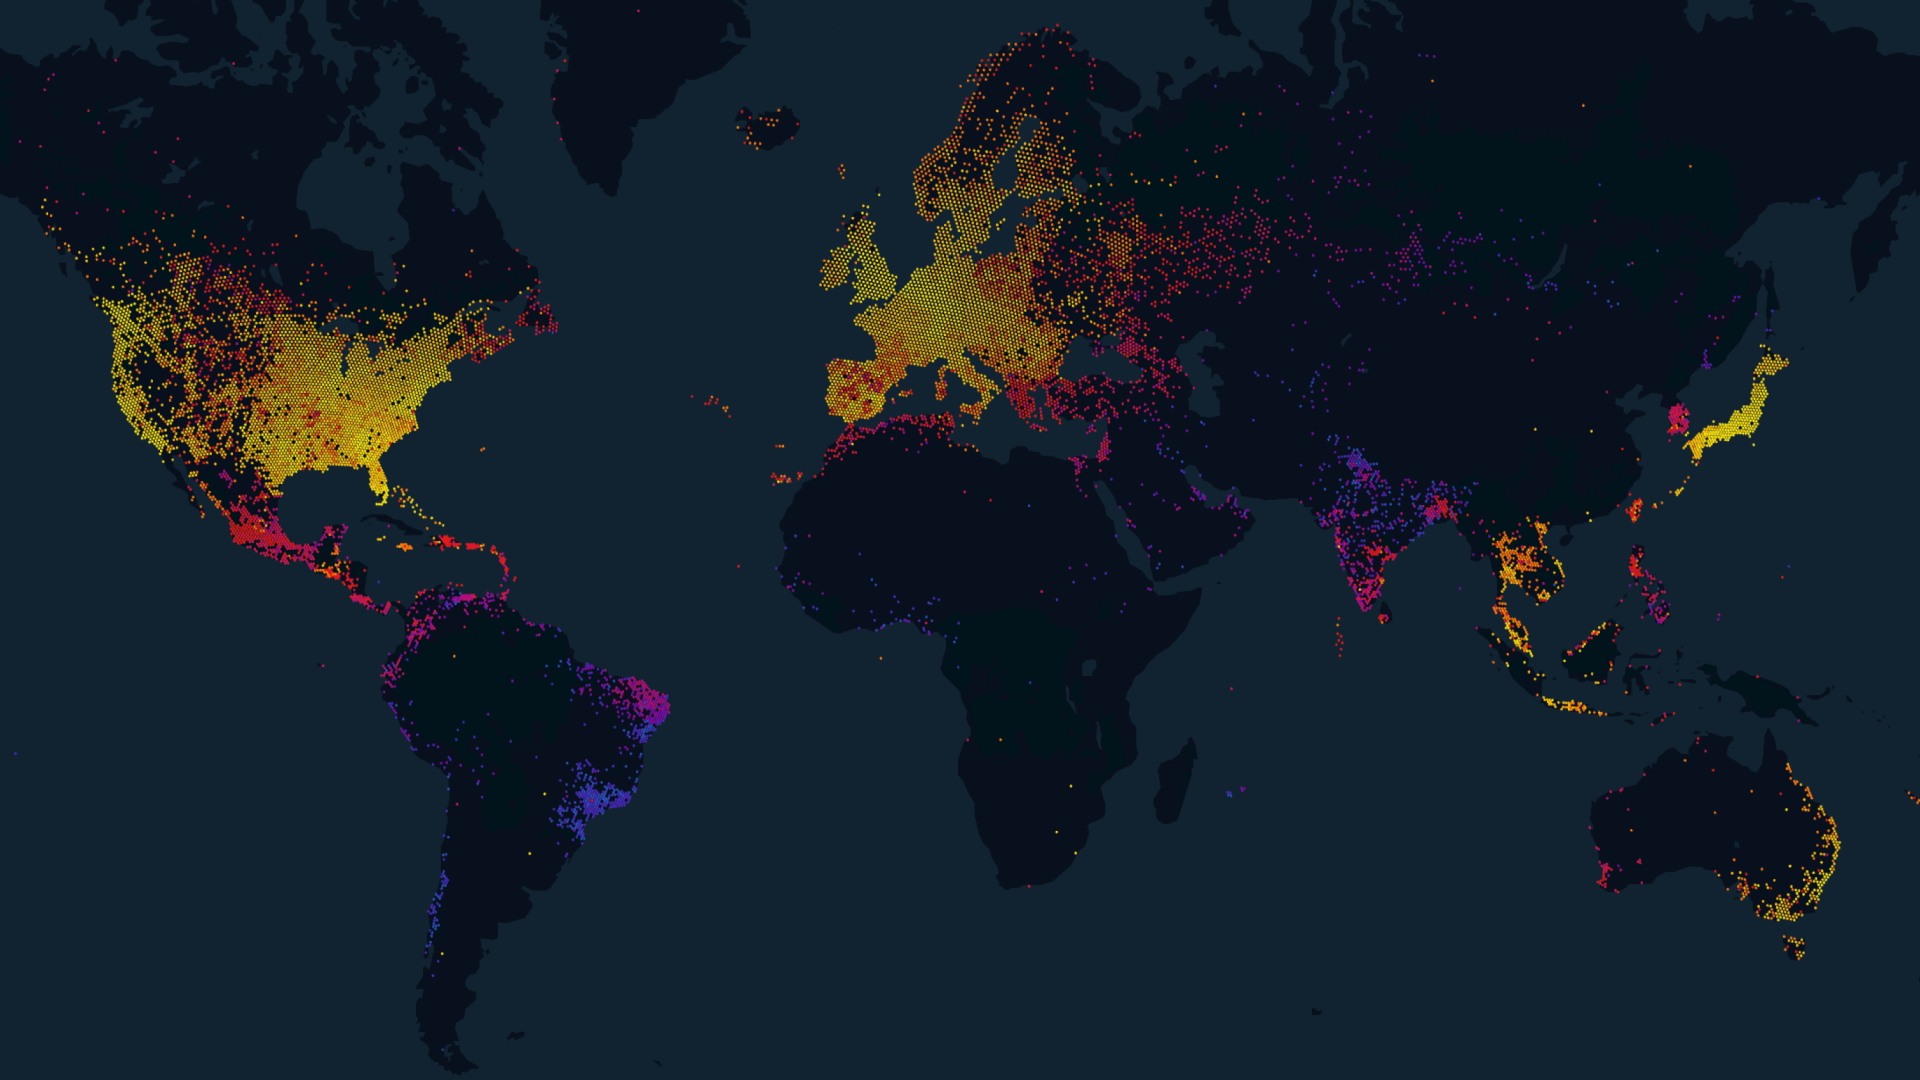
00:00:00 03



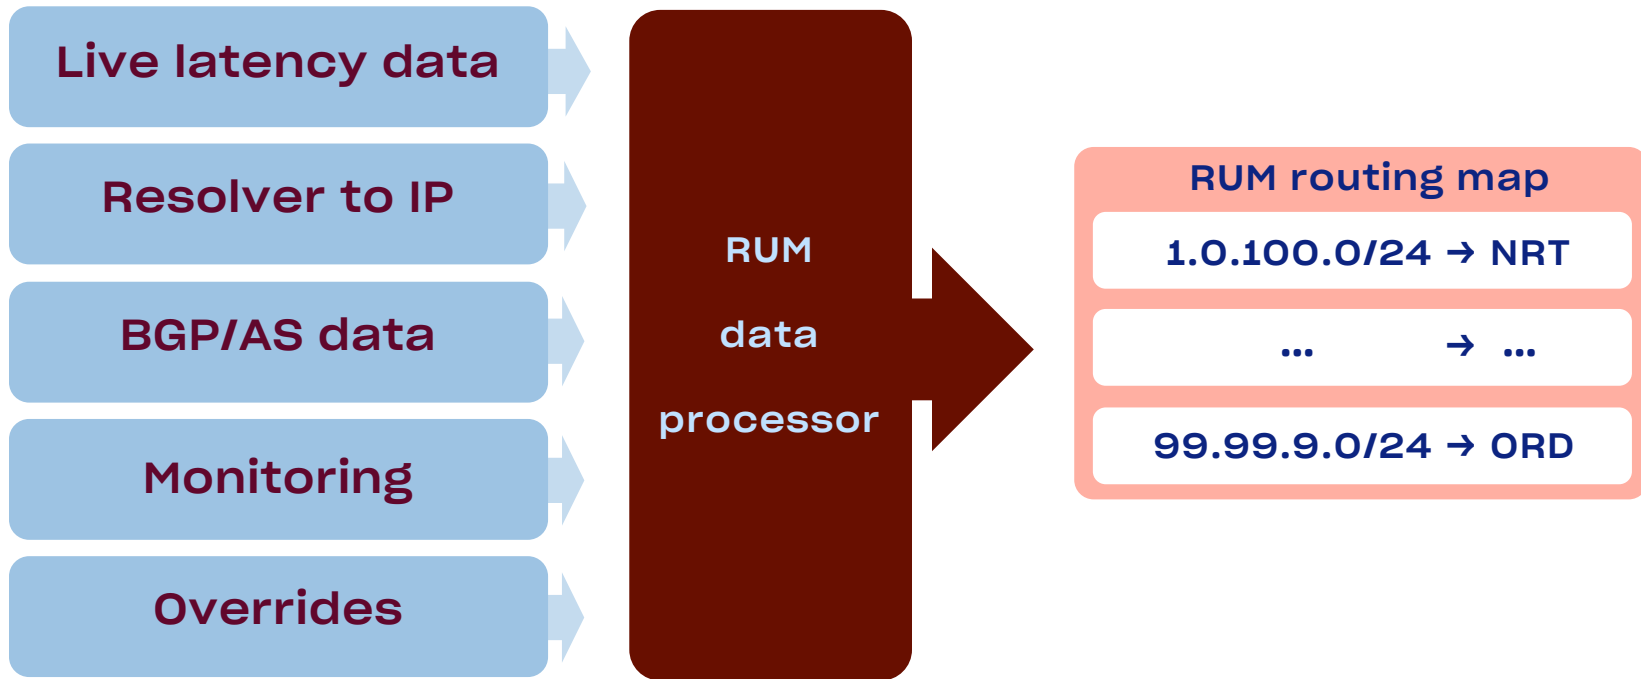
# GSLB: Real user monitoring (RUM)








# GSLB: RUM Data processing



# GSLB: RUM-DNS performance



# GSLB: RUM




**Best routing decisions**  
**Full control over traffic**




**Infrastructure required**  
**Complicated data processing**

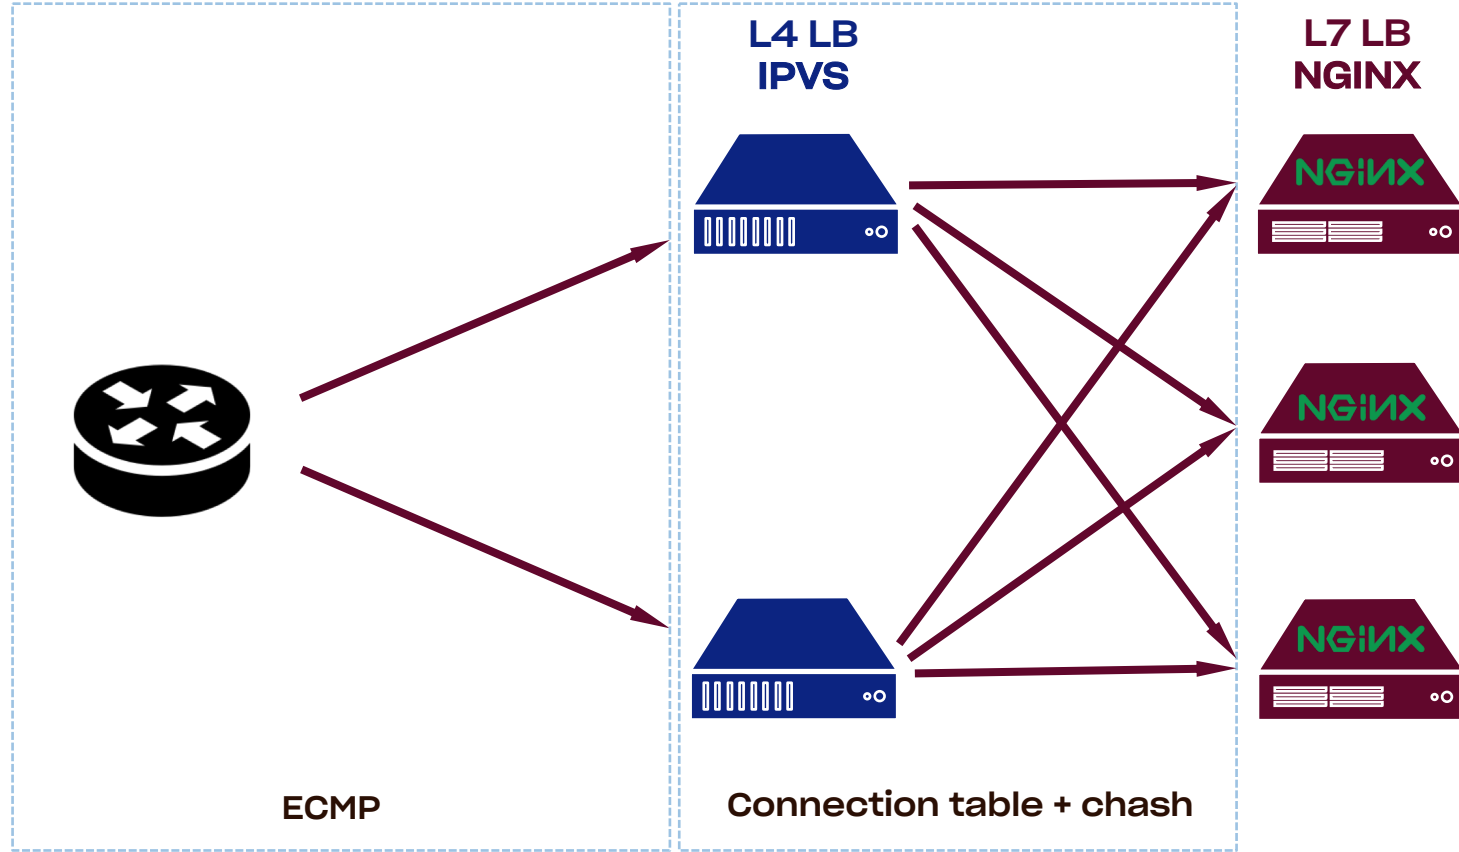
# Inside the Point of Presence




# PoP network architecture



# PoP software architecture

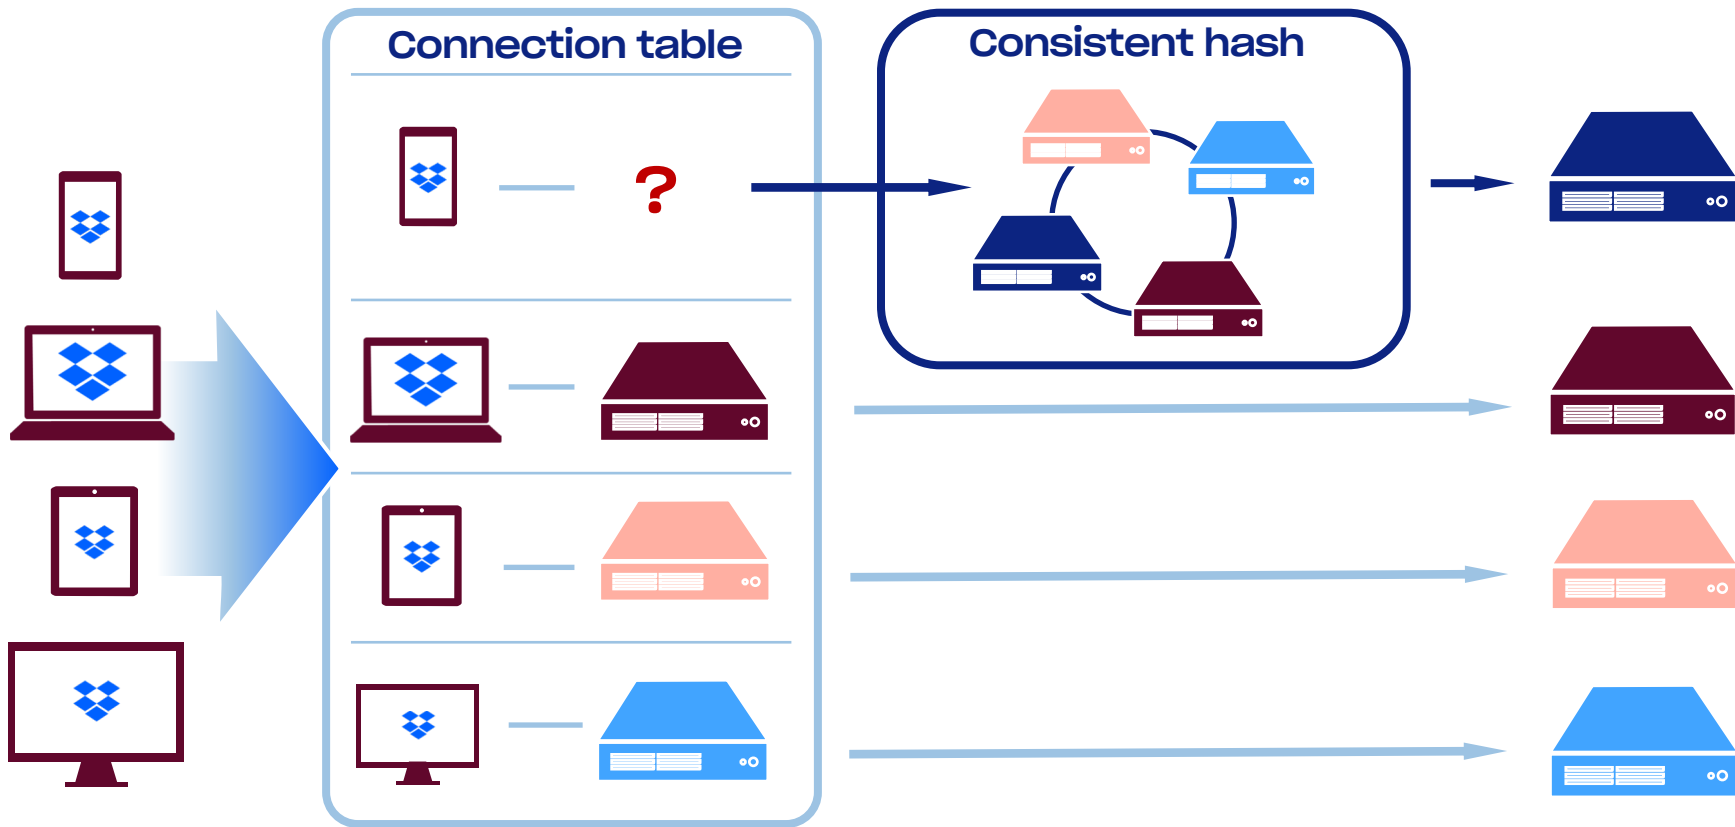


# PoP software architecture

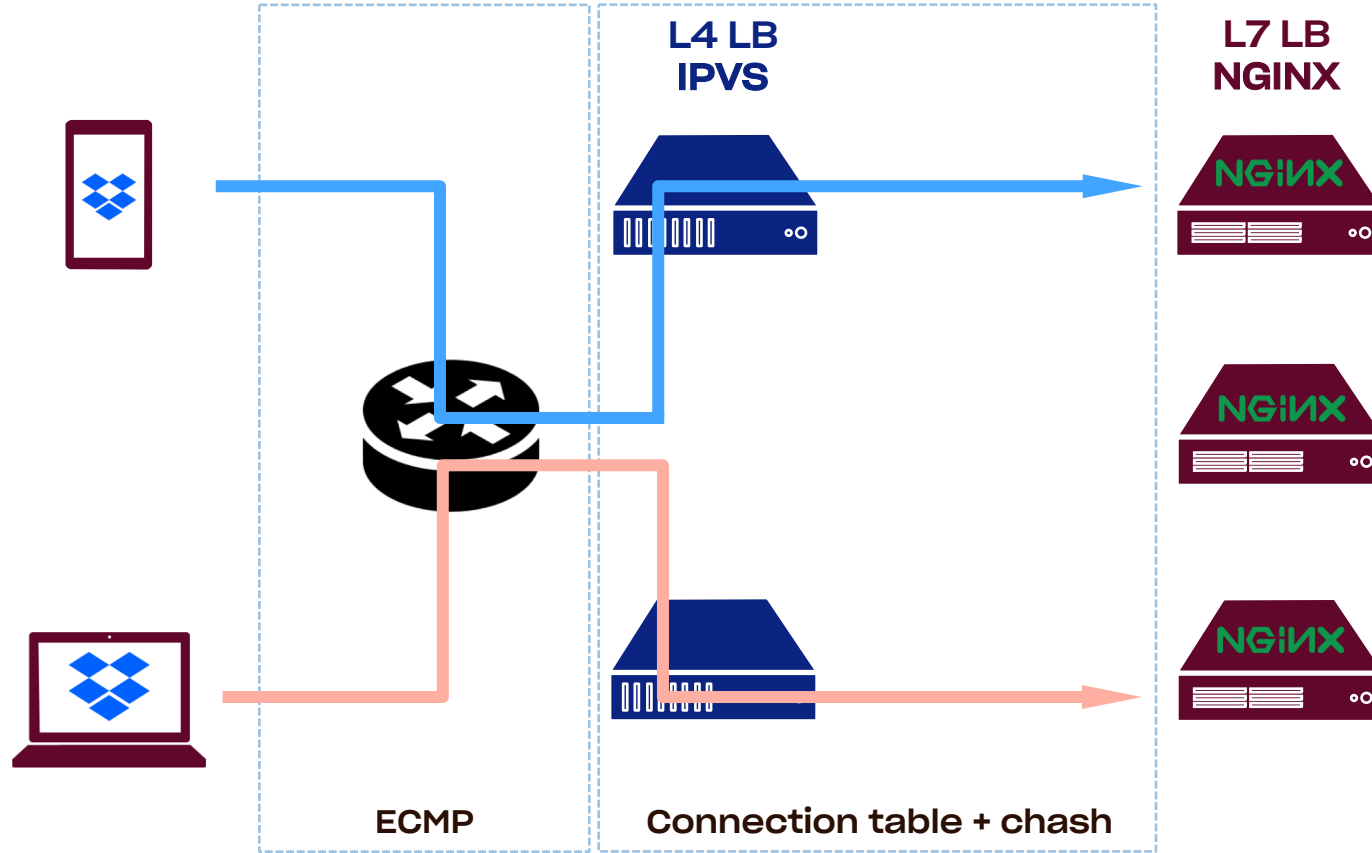




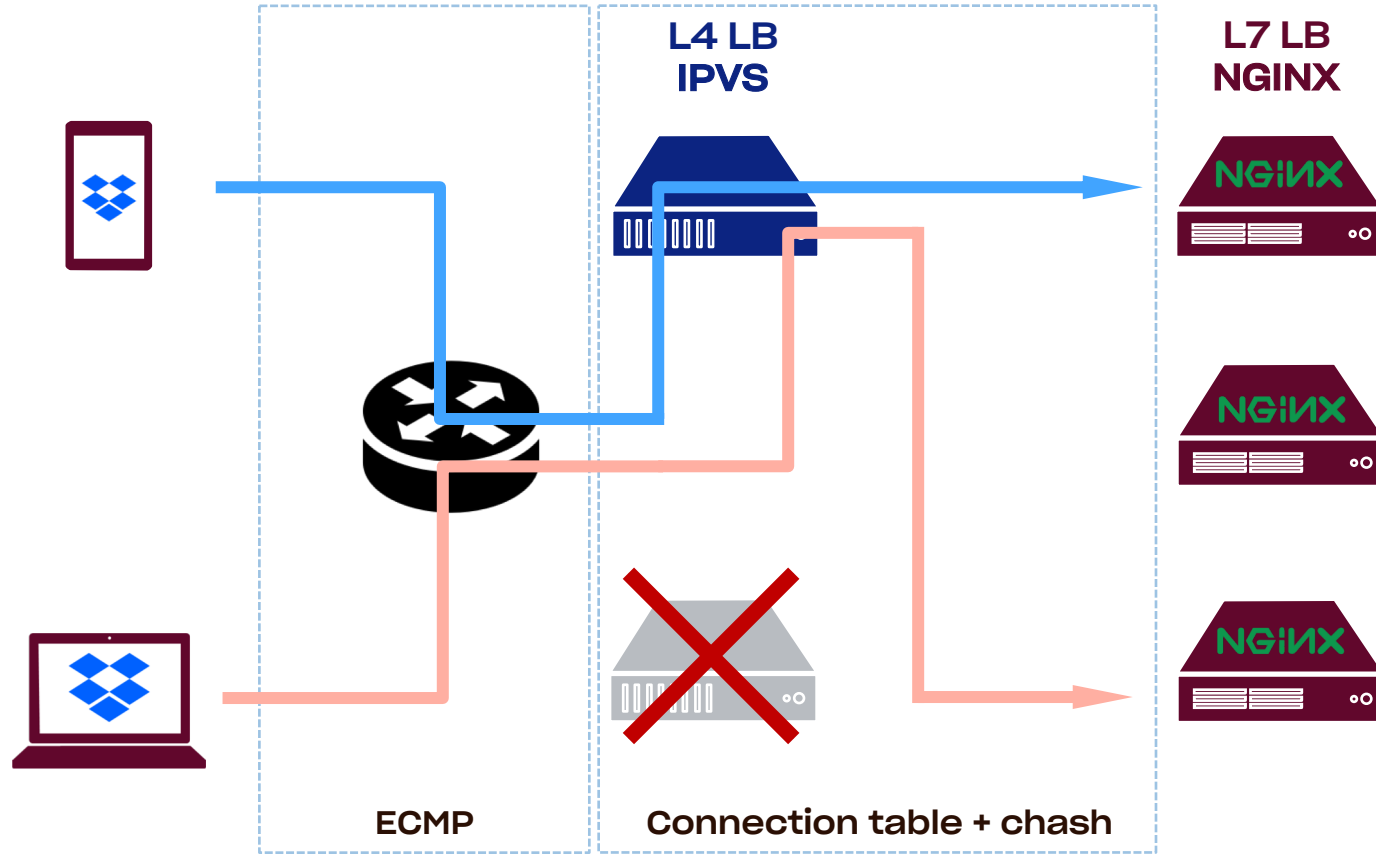
# PoP software architecture: IPVS



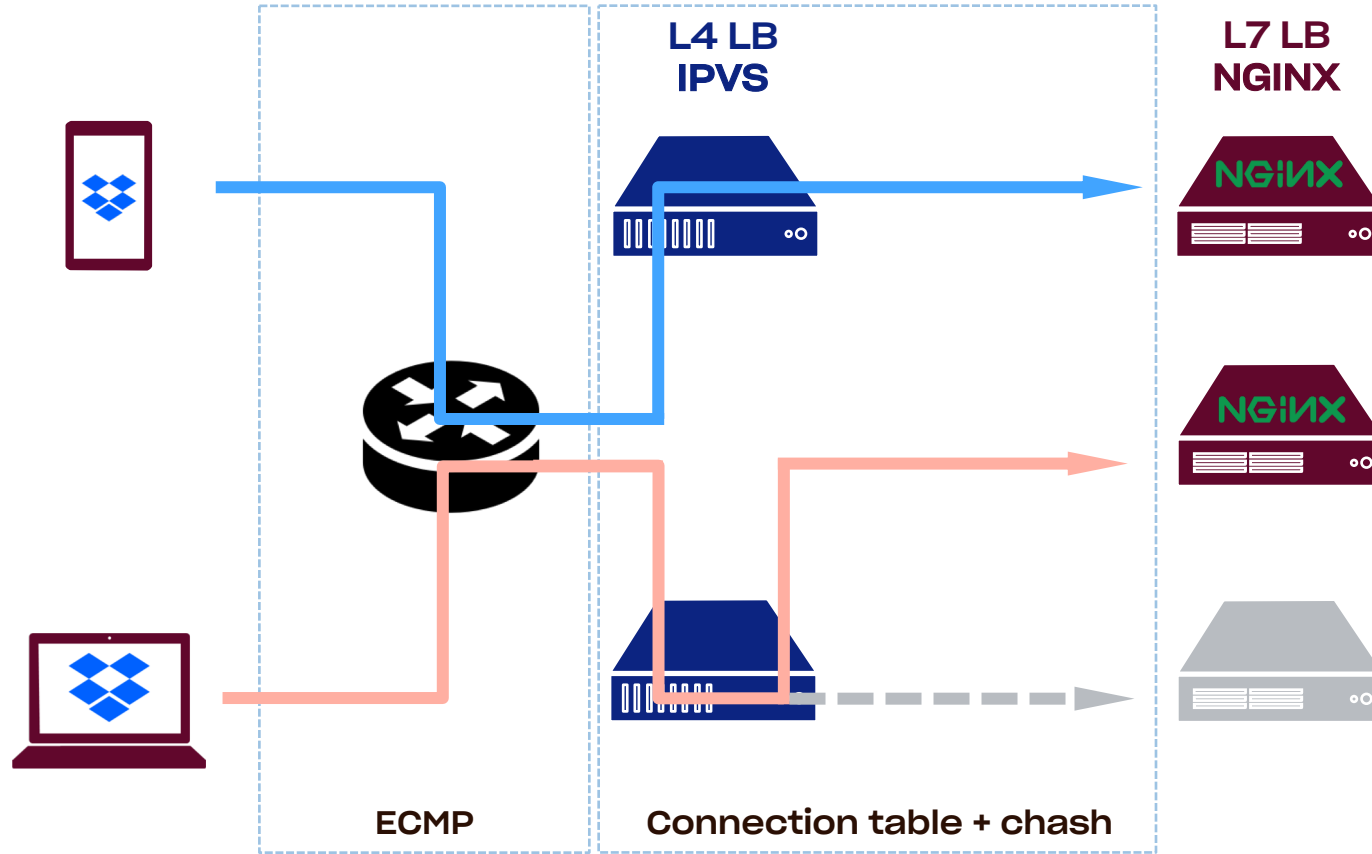
# PoP software architecture




# PoP software architecture: failure modes




# PoP software architecture: failure modes




# NGINX lifecycle




# UMS



# TCP: Fair queueing



# TCP: Pacing





# Summary


Dropbox Edge network

Why latency matters

GSLB approaches

PoP architecture

# We are hiring!





fb: oleg.guba

email: oleg@dropbox.com